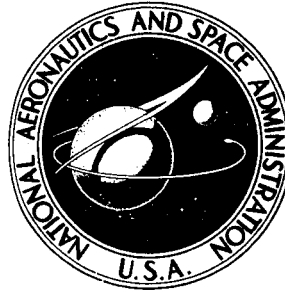


**NASA CONTRACTOR
REPORT**



NASA CR-2460

NASA CR-2460

**A SYSTEM TO GEOMETRICALLY RECTIFY
AND MAP AIRBORNE SCANNER IMAGERY
AND TO ESTIMATE GROUND AREA**

by M. M. Spencer, J. M. Wolf, and M. A. Schall

Prepared by

ENVIRONMENTAL RESEARCH INSTITUTE OF MICHIGAN

Ann Arbor, Mich. 48107

for Lyndon B. Johnson Space Center



NATIONAL AERONAUTICS AND SPACE ADMINISTRATION • WASHINGTON, D. C. • OCTOBER 1974

1. Report No. NASA CR-2460		2. Government Accession No.		3. Recipient's Catalog No.	
4. Title and Subtitle A SYSTEM TO GEOMETRICALLY RECTIFY AND MAP AIRBORNE SCANNER IMAGERY AND TO ESTIMATE GROUND AREA				5. Report Date OCTOBER 1974	
				6. Performing Organization Code	
7. Author(s) M. M. Spencer, J. M. Wolf, M. A. Schall				8. Performing Organization Report No. ERIM 190100-28-T	
9. Performing Organization Name and Address Environmental Research Institute of Michigan Infrared and Optics Division P. O. Box 618 Ann Arbor, Mich. 48107				10. Work Unit No. Task IX	
				11. Contract or Grant No. NAS #9-9784	
12. Sponsoring Agency Name and Address National Aeronautics and Space Administration Lyndon B. Johnson Space Center Earth Observations Division Houston, Texas 77058				13. Type of Report and Period Covered Technical Report 1 February 1973 through 31 January 1974	
				14. Sponsoring Agency Code	
15. Supplementary Notes Dr. Andrew Potter/TF3 is Technical Monitor for NASA					
16. Abstract A system of computer programs has been developed which performs geometric rectification and line-by-line mapping of airborne multispectral scanner data to ground coordinates and estimates ground area. As inputs, the system requires aircraft attitude and positional information (furnished by ancillary aircraft equipment); it also needs ground control points. The geometric correction and mapping procedure locates the scan lines, or the pixels on each line, in terms of map grid coordinates. The area estimation procedure gives ground area for each pixel or for a predesignated parcel specified in map grid coordinates. The results of exercising the system with simulated data showed the uncorrected video and corrected imagery and produced area estimates accurate to better than 99.7%. Real input data, though prepared, have not yet been processed by the system.					
17. Key Words Geometric correction Scanner imagery Imaging techniques Map fitting				18. Distribution Statement STAR Subject Category: 13	
19. Security Classif. (of this report) UNCLASSIFIED		20. Security Classif. (of this page) UNCLASSIFIED		21. No. of Pages 47	
				22. Price \$4.00	

NOTICES

Sponsorship. The work reported herein was conducted by the Environmental Research Institute of Michigan for the National Aeronautics and Space Administration, Johnson Space Center, Houston, Texas 77058, under Contract NAS 9-9784, Task IX. Dr. Andrew Potter/TF3 is Technical Monitor. Contracts and grants to the Institute for the support of sponsored research are administered through the Office of Contracts Administration.

Disclaimers. This report was prepared as an account of Government-sponsored work. Neither the United States, nor the National Aeronautics and Space Administration (NASA), nor any person acting on behalf of NASA:

- (A) Makes any warranty or representation, expressed or implied, with respect to the accuracy, completeness, or usefulness of the information contained in this report, or that the use of any information, apparatus, method, or process disclosed in this report may not infringe privately owned rights; or
- (B) Assumes any liabilities with respect to the use of, or for damages resulting from the use of any information, apparatus, method, or process disclosed in this report.

As used above, "person acting on behalf of NASA" includes any employee or contractor of NASA, or employee of such contractor, to the extent that such employee or contractor of NASA or employee of such contractor prepares, disseminates, or provides access to any information pursuant to his employment or contract with NASA, or his employment with such contractor.

Availability Notice. Requests for copies of this report should be referred to:

National Aeronautics and Space Administration
Scientific and Technical Information Facility
P. O. Box 33
College Park, Maryland 20740

Final Disposition. After this document has served its purpose, it may be destroyed. Please do not return it to the Environmental Research Institute of Michigan.

PREFACE

This report describes part of a comprehensive and continuing program of research in multispectral remote sensing of the environment from aircraft and satellites. The research is being carried out for NASA's Lyndon B. Johnson Space Center, Houston, Texas, by the Environmental Research Institute of Michigan (formerly the Willow Run Laboratories, a unit of The University of Michigan's Institute of Science and Technology). The basic objective of this program is to develop remote sensing as a practical tool for obtaining extensive environmental information quickly and economically.

One pressing need identified during the past five years of working with government, NASA, and university users of remote sensing data, is that of generating scanner data in a form directly comparable with maps or rectified photos. Recent emphasis on area measurement for assessment of pond areas and perimeters in waterfowl management on the Great Plains and for the proposed ASCS (Agricultural Stabilization and Conservation Service) acreage compliance program demands that a more unified look must be taken at the correction of scanner data distortions caused by scanning geometry and aircraft platform instability.

Results of an initial study of the airborne scanner cartography problem are reported herein. This study identifies the quantitative effects of aircraft instabilities and scanner geometry on the scanner data, specifies correction techniques utilizing information from ancillary sensors and ground truth data, and develops an area algorithm. The results of this study led to the recent development of a set of computer programs written in FORTRAN IV and comprising the Scanner Imagery Correction and Line-by-line Orientation Programs (SICLOPS) System.

This report describes:

- (1) The SICLOPS System, which performs line-by-line mapping of airborne scanner data to ground coordinates and estimates ground areas
- (2) A simulator for generating input data with which to test the SICLOPS System and for use as a tool in system error analysis
- (3) System operation with simulated data to produce simulated video and a corrected image, as well as to verify ground measurements obtained with the area algorithm
- (4) System operation with data containing simulated errors to evaluate the fitting techniques used
- (5) Preparation of real input data for future processing in the SICLOPS System

The research covered in this report was performed under Contract NAS 9-9784, Task IX, and covers the period from 1 February 1973 through 31 October 1973. Dr. Andrew Potter has been Technical Monitor for NASA. The program was directed by R. R. Legault, a Vice-President of the Environmental Research Institute of Michigan (ERIM); J. D. Erickson, Principal Investigator and Head of the ERIM Information Systems and Analysis Department; and R. F. Nalepka, Head of the ERIM Multispectral Analysis Section. The ERIM number for this report is 190100-28-T.

The authors acknowledge the administrative direction provided by Mr. R. R. Legault, Dr. J. D. Erickson, and Mr. R. F. Nalepka; the technical counsel furnished by Messrs. R. J. Kauth, W. A. Malila, and I. J. Sattinger; the editorial support of B. M. Sanderson; and the secretarial services of Miss D. Dickerson.

The authors also appreciate the cooperation of Dr. Carl J. Born of the Defense Mapping Agency Topographic Center in sending the Casa Grande (Arizona) Test Site Data, and the support given by the following people from NASA/JSC: Mr. Richard H. Baker (Aircraft Project Manager for this task's flight mission), Mr. Charles Davis, Mr. John Dietrich, Mr. R. Ligons, and Mr. R. Kyle. The technical counsel furnished by Mr. Kraft and Mr. Diesel of the Litton Corporation is likewise appreciated.

CONTENTS

1. Summary	7
2. Introduction	8
3. Description of the SICLOPS System	15
3.1 Preparation of MSS Data	17
3.2 Preparation of ADAS Data	17
3.3 Preparation of LTN-51 Data	19
3.4 Preparation of Ground Truth Data	22
3.5 Preliminary Geometric Correction and Fitting	23
3.6 Final Data Processing	28
4. System Demonstration Using Simulated Data	32
4.1 ADAS Simulator	32
4.2 LTN-51 Simulator	32
4.3 Ground Truth Data Simulator	33
4.4 Video Simulator	33
4.5 Preliminary Geometric Correction and Fitting	35
4.6 Final Data Processing	36
4.7 Results	37
5. Comments on Applying the SICLOPS System to Real Data	41
6. Conclusions and Recommendations	42
Bibliography	44
Appendix: Axis Systems	45
Distribution List	46

FIGURES

1. Data Flow Diagram of the NC130B System	13
2. Functional Flow Diagram of SICLOPS Systems	16
3. Preparation of MSS Data	18
4. Preparation of ADAS Data	18
5. Preparation of LTN-51 Data	20
6. Allocation of Pixel Area Using Area Algorithm	30
7. Simulated Video	38
8. Corrected Image	38

TABLES

1. Area Measurement of "Test" Pattern on Run 1	40
2. Area Measurement of "Test" Pattern on Run 2	40

A SYSTEM TO GEOMETRICALLY RECTIFY AND MAP AIRBORNE SCANNER IMAGERY AND TO ESTIMATE GROUND AREA

SUMMARY

This report describes a system of computer programs (SICLOPS)* which has been developed to perform geometrically corrected line-by-line mapping of airborne multispectral scanner data in ground coordinates and to estimate ground areas. As inputs, the system requires aircraft attitude and positional information furnished by ancillary aircraft equipment, plus ground truth information giving the location, in map grid coordinates, of a number of recognizable landmarks on the ground.

The effects of aircraft motion are first partially compensated by applying the navigational data to the data points corresponding to the landmarks in the scanner imagery. After this preliminary geometric correction, the correct data points and the ground truth information are used in a two-stage fitting process. The first fit is special, designed to eliminate errors not handled by the second, which is a standard cartographic fit. Final data processing does the actual mapping; applied to all data points, it combines the preliminary geometric correction and the fitting corrections into one operation, converting line or pixel locations into map grid coordinates. To obtain ground area values for individual pixels or for predesignated parcels specified in graymap coordinates, the area algorithm (area estimation procedure) is used in the data processor.

The system was exercised with simulated data to demonstrate its feasibility. Results show a dramatic improvement in the corrected over the uncorrected imagery. Area measurements were also made on the simulated data which included simulated errors; these errors are round-off and quantization errors (which are characteristic of the real input data) and systematic errors. In the presence of significant simulated errors, the area estimates were accurate to better than 99.7%.

Real input data have been prepared for the system but, pending generation of ground truth information, have not as yet been used.

*Scanner Imagery Correction and Line-by-line Orientation ProgramS.

INTRODUCTION

Airborne scanner imagery is typically subject to distortions. In the scanner art, interest has traditionally focussed on special properties unique to the scanner image. The special aspects of, say, a thermal or perhaps a multispectral image have captured attention, while geometric fidelity, though technically possible, has been neglected.

Scanners have meanwhile emerged from the demonstration stage, and are being placed in decidedly expensive environments by serious users who expect more than a demonstration to result. Extensive areas are being covered, and automatic means of recognition are available. The inability to locate objects or measure either distance or area from scanner data is now becoming an important limitation. Questions are being asked such as: What is the acreage of the crop detected? The location of the trees detected as being diseased? The acreage of a crop within given boundaries?

Ultimately, in order to meet this new range of requirements, it will be necessary to produce a rectified and faithfully scaled image from scanner data. The problems implicit in this task are different from those associated with aircraft photography or with spacecraft scanner data. In photography, two-dimensional distortions caused by aircraft attitude perturbations must be corrected for, but need be made only once per frame. In scanner data, corrections for distortion are also two-dimensional, but aircraft attitude corrections are needed much more frequently. In comparison to spacecraft, aircraft are much less stable as sensor platforms, especially if flown at low altitude in turbulent air.

We need to think carefully about what it is we are attempting to achieve. A facile comparison to photography can lead us astray. Photographs (taken with a nearly vertical camera) can be rectified by linear transformations performed on the image as a whole. A differential stretching so attained will fit the scaled image to the ground. This method cannot be successful with scanner images for the following reasons: (1) aircraft attitude may change during the sequence of scanner observations that form the scanner image, and (2) the scanner image has scanner angle as one coordinate, rather than scaled ground distance. The usual video scanner image is just a by-product, not a necessary output of the system. Our objective should not be to transform this video image, but rather to operate on the total set of video signals and other ancillary data to produce rectified imagery. This has not been done previously .

At the present time, no operational use is being made of information from the LITTON inertial navigation system on board NASA's NC130B aircraft to correct motion-related distortions in MSS data. The same holds true for the limited amount of aircraft platform orientation information available in ERIM's C-46. No ancillary inertial navigation systems are available in ERIM's C-47.

Approximate rectification has been done in the past, primarily to remove gross distortions distressing to the eye. Roll correction is common, as is the adjustment of V over H in the imaging system. Sometimes approximate yaw is removed, as well as the tangent (panoramic) distortion. But these are essentially cosmetic corrections. For example, the fluctuating component of roll may be removed dynamically, but a constant roll error may remain. Or the average yaw may be removed, but not the fluctuating component, etc. Such cosmetic measures clean up the picture so that it becomes more pleasing to the eye, but accurate location of ground points by measurement is still not feasible.

Before actually producing rectified imagery, one must consider the problem of associating with each pixel a ground location in terms of some set of grid coordinates. To locate that ground point corresponding to a particular element of a scanner picture of relatively flat terrain, eight quantities are required as a function of time: the location of the aircraft in x, y and in altitude above the ground, three angular quantities describing the direction in which the scanner is pointing, and also two orthogonal components of the velocity vector.

The best use of the angular data would probably be to correct the attitude of the scanner during flight. This would have the effect of reducing the entropy of the data as recorded. The ground would be scanned in a more systematic manner, and it would then be simpler to find which picture element corresponds to an assigned ground point — an important matter quite different than finding the ground location of a picture element.

The servo adjustment of a scanner in angle seems at first thought uninviting, inasmuch as scanners are heavy and have a large gyroscopic moment. However, one does not want to rotate the scanner at all; rather, it is the aircraft that should be allowed to move, while the scanner should remain stable. To accomplish this, the scanner can be mounted in gimbals and the servo need only overcome bearing friction. The present study does not deal with this problem. Instead, it addresses the problem of correcting scanner data obtained from a conventionally mounted unit.

The angular data, one would hope, might be adequate, although there may be calibration errors and adjustments necessary to bring the north reference of the system into registry with an arbitrary local grid north. The terrain clearance is also potentially accurate enough. On the other hand, the absolute position of the aircraft cannot be expected to be sufficiently accurate. The drift of inertial navigators is impressively low, on the order of a mile an hour. (As a sidelight, Columbus on his first voyage was out in his navigation by just one mile per hour. It is generally believed that he falsified his data, as no human navigator of that time could have done so poorly.)

The accuracy goal in the present study is, implicitly, to locate the resolution patch to an accuracy of about the dimension of that patch. This is on the order of six meters for the MSS

at an altitude of 3000 meters. In a 1-hour flight the accumulated navigational error may be a mile, or on the order of 1600 meters. Clearly, it will be necessary to use at least one landmark to correct the position—and in fact more than one will be needed, for the additional error accumulated on a 5-minute data run would be about 80 meters. True, one might employ a ground-based navigational system and avoid using landmarks, but in principle such systems simply provide landmarks of an electronic nature.

The requirement for landmarks precludes airborne processing. One cannot know the correction until the whole of the data is collected. One is thus committed to post-flight processing, which for most purposes is not a handicap.

A question of calibration arises with any system of instrumentation. In this case the problem seems particularly acute. Data from three sources are being combined in a new way. The approach adopted is to make the system self-calibrating, through extension of the landmark data set to include a statistical sample of ground-truth locations. In principle, it then is possible to remove the errors during post-flight processing.

Assembling this ground truth data, admittedly a tedious process, may not be required in other than a calibration flight. In any case such a set of data is essential to learn whether a correction system is operating properly.

The system of computer programs* developed in this study comprises the Scanner Imagery Correction and Line-by-line Orientation ProgramS (SICLOPS) System. The correction process it employs is a least-squares fitting, but not the usual fitting used in photogrammetric adjustment. Scanner pictures, even after approximate tangent correction, may contain nonlinear errors which cannot be removed by a linear 2×2 correction matrix. The most probable cause of such errors is in the roll sensing and scanner rotational synchronization provisions which together determine the "scanner angle."

The correction regime advocated here first prepares unfitted positional estimates for the set of ground truth points, and then prepares an estimate of an additional (temporary) quantity, which is in essence the distance that the point appears to lie to the side of the aircraft. Next, an intermediate 2×4 matrix is fitted, which is used only for the estimation of the error in the scanner angle. One element of this matrix is interpreted to give an estimate of the scanner angle error, and the scanner angle is corrected by this amount. New estimates are then made of the positions of the check points. These estimates hopefully differ from the initial estimates in that the remaining errors are linear, so a 2×2 matrix can correct them. A 2×2 matrix is then fitted and used in subsequent correction of the data to give a fitted set of estimates from

*Written in FORTRAN IV for an IBM 360 computer.

which the rms errors can be evaluated and used as the best available test of the accuracy of the overall process.

These same corrections may then be applied to the scanner data in general. From the fitting process one obtains: a knowledge of an x and y translation to be applied to the data, a scanner angle correction, and a 2×2 correction matrix. It should be emphasized that although the correction process described requires a somewhat lengthy program, this fitting program is applied only to the quite small set of ground truth values, not to the voluminous scanner data. The computing time required is thus modest.

The location of ground points is essential as a preliminary to the measurement of ground areas. The actual area inspected in a single resolution element is irregularly shaped, and may overlap neighboring resolution elements in varying degrees. Such elements cannot be summed to give the area of a ground region. The approach taken has been to dissect the image, conceptually, into zones about the center location of each resolution element, giving each such zone a boundary distinctly separating it from adjacent ones. These zones are in one-to-one correspondence with the resolution elements, and have the property that their individual areas may be summed to give the correct area within some larger boundary.

A somewhat difficult problem must be faced in making use of the area data so gained. An answer such as "321 acres of wheat" is not useful. One must be able to state a boundary within which that 321 acres of wheat was found. It is not a simple matter to decide by mathematical means that a given point is within a given boundary. The matter of "inside" versus "outside" is a human perceptual concept, which takes on some complexity when treated mathematically.

Present system programs do not carry the problem through the phase of providing this necessary boundary, but the method adopted was chosen with this ultimate application in mind. The points are treated as elements of scanner lines projected onto the ground. As a simplification, these are treated as straight lines; the attitude of the scanner is presumed to remain unchanged during any one scan line. It then becomes feasible to locate the intersection of the scan line with the boundary of a simple polygon such as a triangle or a convex quadrilateral. An arbitrary boundary can be approximated as a sum of such simple figures. If the scan line enters a convex polygon, there will in general be only one point of entry and only one of exit. Image points between the point of entry and the point of exit are "inside," others are "outside,"

The treatment in terms of scan lines permits a compactness of data. One may simply transmit the parameters of a given scan line and the rule for finding the ground coordinates for image elements on that line, rather than transmitting the many hundreds of coordinate pairs locating each image element. In addition, an efficient area algorithm may be based on integration in strips along a scan line, rather than on zone-by-zone integration. Many adjacent

image elements typically have the same "recognition," Hence, one need only compute the ends of such a coherent strip—which means large savings in computer effort.

This brings us through the full process, yet not quite to the end of the study effort. Still required is confirmatory testing, for which we have prepared a series of programs. Real data taken with the scanner and associated system are not appropriate to the task of confirming that the programs function as intended. Here, by analogy, one has constructed an analytical balance. It is not appropriate to test this balance by using objects of uncertain or unknown weight. Instead one wants known standard weights.

In the present study we programmed a series of simulators which generate synthetic data representing outputs from a scanner system, ancillary systems, and ground truth. Into two of these simulators (the ADAS and LTN-51) we introduced errors corresponding to the round-off and quantization errors of the real data and, in addition, inserted systematic errors as well. The data from all the simulators are, in effect, the standards against which the new programs may be tried to confirm that they indeed function (apart from any instrument-related shortcomings which may be present in the real data). Real data were prepared during the present study but have not yet been processed by the SICLOPS System.

We elected to convert all data to MKS, with angles in radians, and to use the metric UTM grid in expressing ground locations. The data have various units: knots, degrees, approximate milliradians, and feet. It would be desirable to use the land measure grid, yet there are three such grids used in the U.S.A., and all are in fact irregular.

The source of real data for this study was NASA's NC130B system, which is an elaborate recording system on the NC130B aircraft. The airborne equipments of interest to this study were the 24-channel Multispectral Scanner (MSS), the Litton navigation system, and the radar altimeter. The master clock, IRIG-A, which furnished reference time to all recordings made on the aircraft, was an important adjunct of the NC130B system for this study.

Figure 1 is a data flow diagram of this NC130B system showing the 24-channel MSS and ancillary equipment. There are two 14-track AR-1600 magnetic tape recorders aboard the aircraft. The tape recorder shown on the right side of the diagram records the 24-channel scanner data on 12 of its tracks. The tape recorder on the left side of the diagram records the outputs of ancillary equipment on three tracks: IRIG-A, the LTN-51 record which is a BCD output of the Litton navigation system, and the ADAS (Airborne Data Annotated System) record.

The ADAS record contains radar altimeter and the Litton platform information. (A second ADAS record is also recorded along with the Scanner Housekeeping information on track 13 of the tape recorder on the right-hand side of Fig. 1; this second ADAS record is unsynchronized with scanner data and is not to be confused with the ADAS record referred to in the rest of the

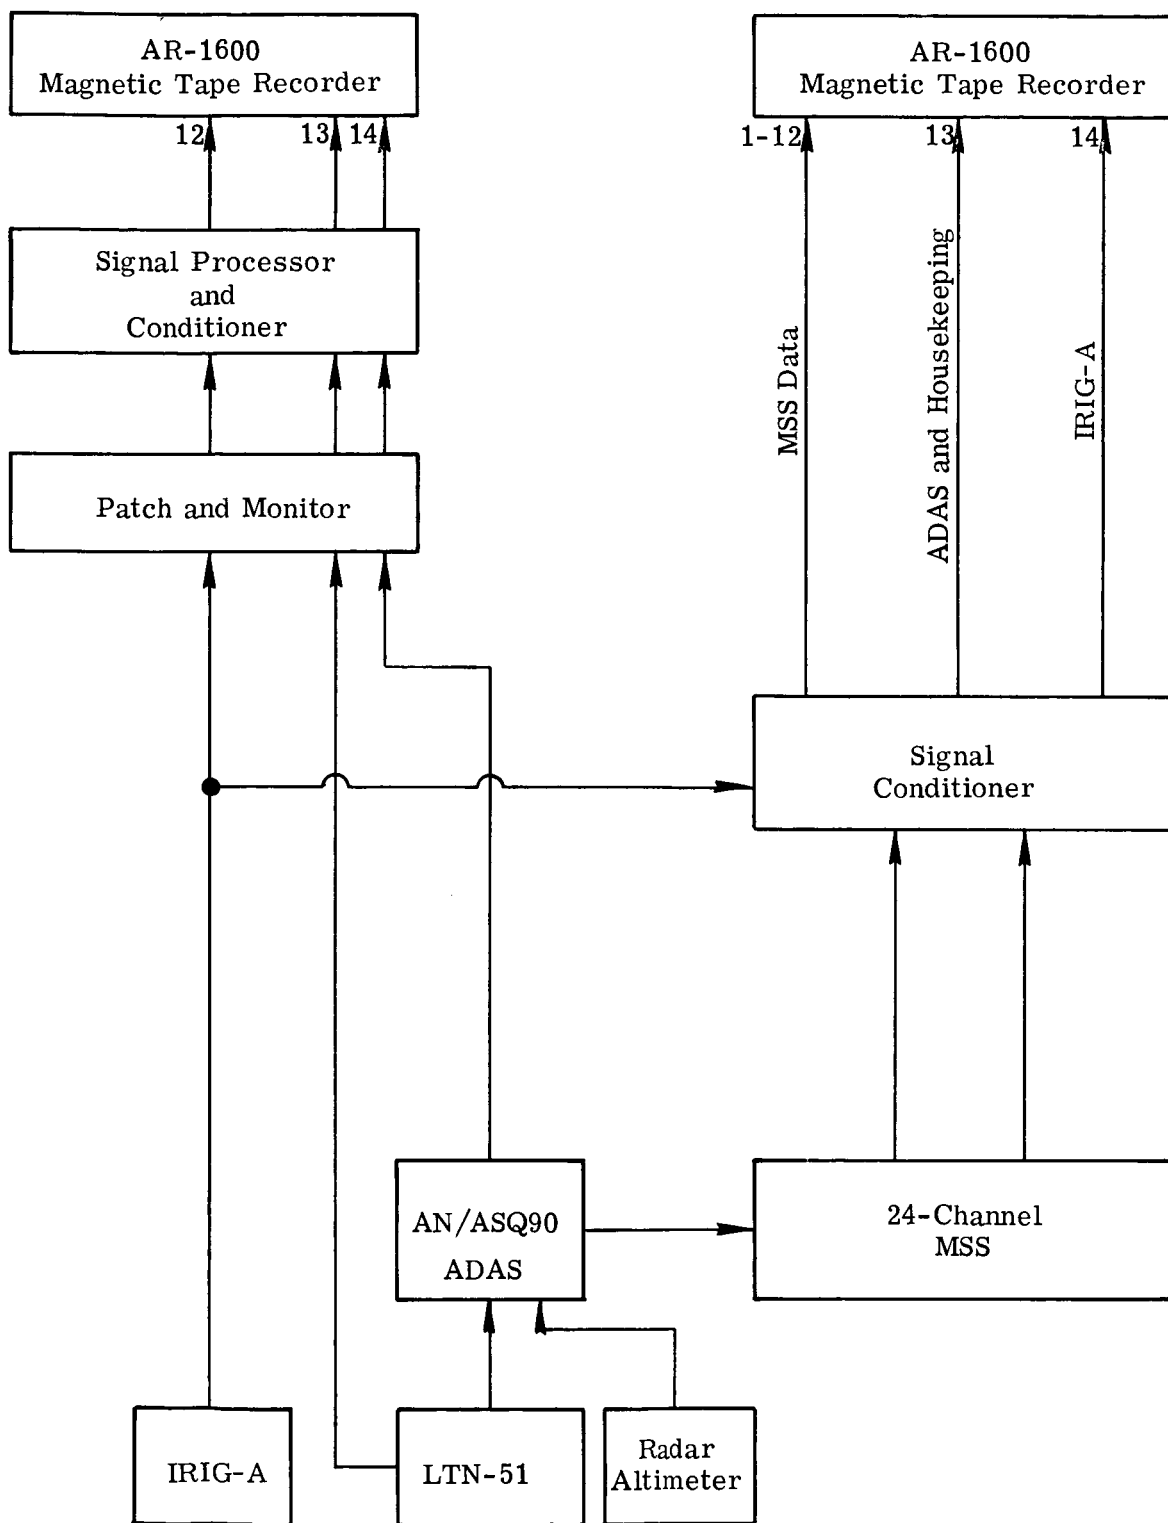


FIGURE 1. DATA FLOW DIAGRAM OF THE NC130B SYSTEM

report.) The IRIG-A is recorded on the 14th track of the tape recorder on the right side of the figure. The data from the LTN-51, the ADAS, and the MSS—as recorded by the NC130B system and converted to computer-compatible tapes—provide the necessary aircraft input data to the SICLOPS System.

The programs of the SICLOPS System presently devoted to data preparation are specially designed to accept this particular input of scanner and aircraft data from the NC130B system. The main programs of the system, however, apply to any imaging airborne scanner system and can be utilized as long as the necessary input data are available.

The SICLOPS System is described in Section 3, along with a description of the data preparation needed for the aircraft input data, the preliminary correction and fitting procedures, and the final data processing which involves the area algorithm. Section 4 describes use of simulated data by the SICLOPS System and the results obtained. Since the system was not exercised with real data, Section 5 presents some thoughts on the subject of possible results one might expect when real data is applied to the system. Section 6 presents the conclusions and recommendations resulting from this study.

DESCRIPTION OF THE SICLOPS SYSTEM

A functional flow diagram of ERIM's SICLOPS System appears in Fig. 2. There are three kinds of real aircraft data: (1) the MSS record containing video data from the 24-channel multi-spectral scanner, (2) the ADAS record (Airborne Data Annotation System), and (3) the LTN-51 record containing information from the Litton navigation system. The other data input is ground truth data, making four inputs in all.

The real aircraft input data is sent to the Aircraft Data Preparation block. The MSS video data is reformatted and the scan line numbers and their corresponding times are stripped from the record for use. The ADAS and LTN-51 data needed by the system are removed from the records and repaired if there are time gaps in the data or if wild values appear. The data are then smoothed, time-corrected, and interpolated. The essential data from the three sources then merge as time-synchronized data. A subsidiary output of this block is a graymap used in the Ground Truth Data Preparation block.

The time-synchronized data are used in the Preliminary Geometric Correction block. The preliminary correction, using the aircraft attitude and positional information, positions pixels within an arbitrary ground coordinate system. This correction is only made on a subset of the data points.

The preliminary corrected data points are provided to the Data Fitting block, where the ground truth points are compared to these points and a fit is made. The fitting procedure is performed in two steps. The first step is designed to eliminate scanner angle error. Adjustment is made to the scanner angle zero and corrected data points corresponding to the ground truth points are again computed. A standard cartographic fit is then made. This fit produces a set of correction parameters to be used in the Data Processor block.

The Data Processor block, using the information from the Data Preparation block, the adjusted preliminary correction, and the standard cartographic fit, processes all the data points. The area algorithm is contained in the Data Processor block. The outputs are either the areas and ground coordinates of each pixel or, alternatively, the area of each predesignated ground parcel. The ground parcel is designated in bounding graymap coordinates.

The complete system has been exercised with simulated input data. The Simulation block (in upper left-hand corner of Fig. 2) contains the four data simulators needed to furnish the system with simulated data. Into two of these simulators, the ADAS and the LTN-51, errors can be introduced corresponding to the round-off and quantization errors of the real data; systematic errors can also be introduced. The data from all the simulators are, in effect, the standards against which the system may be tried, to confirm that it indeed functions apart from any instrumental shortcomings which may be present in the real data.

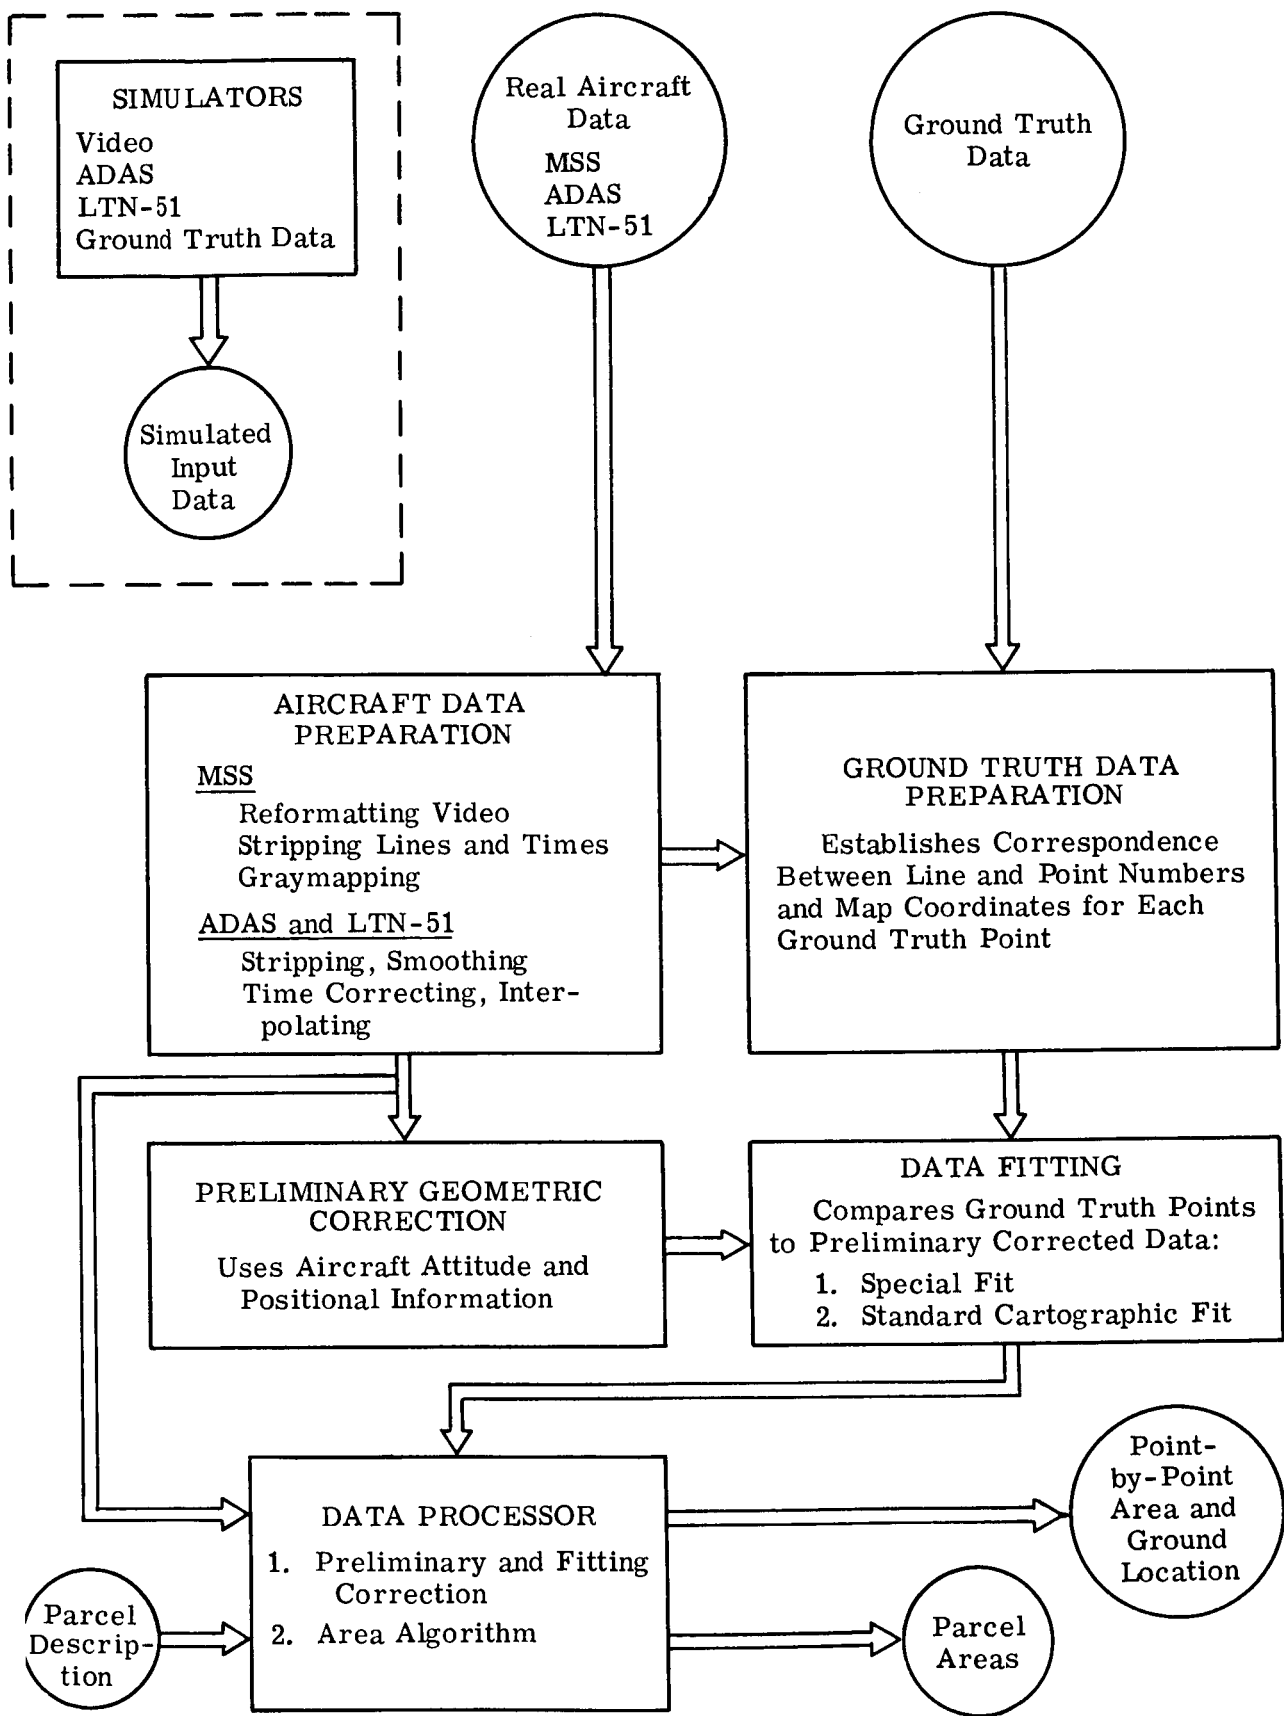


FIGURE 2. FUNCTIONAL FLOW DIAGRAM OF THE SICLOPS SYSTEM

The complete system has not been exercised with real data. A major effort was made by the sponsor in supplying the necessary data, and a major effort was made at ERIM in extracting the desired data from the records supplied, and in inspecting these records for gaps or irregularities in the data. Some record gaps were found, and these have been repaired by linear interpolation. Some wild data points were also found and repaired. All the aircraft input data have been prepared and a few portions of the data have been graymapped. A full set of records, minus the ground truth data, has now been assembled and awaits processing at the preliminary correction and fitting stage.

The SICLOPS System at present consists of a series of programs which process the data sequentially. This sequential feature was a convenience in program development and also provided an economy in computer usage since large blocks of core storage did not have to be reserved during the entire processing to store intermediate data. The developed programs can now be modified for more simultaneous operation.

3.1 PREPARATION OF MSS DATA

The video record, called MSS Data, contains:

- (a) A scanline serial number
- (b) A scanline time relating to the time of the 20th (or other) picture element ("pixel")
- (c) 700 video amplitudes pertaining to the 700 successive pixels in the scan line. The video amplitudes are quantized in 256 gray levels.

The block diagram for preparation of the MSS data appears in Fig. 3. Two possible formats can be processed, UNIVERSAL or MSDS. The program MSS (CCT in Universal format) reformats the video record for compatibility with an existing ERIM program, GRAYMAP. Program MSS also produces another record containing scanline number and scanline time which is an input record to program TMSS. TMSS extrapolates the MSS scanline time to a TMSS time which is related to the midpoint point of the scan line.

The MSDS program handles the MSS data in the MSDS format. The MSDS data however lack the scanline time. A short program "Time Line," which from a few points read from a video print of data computes the slope and intercept of a fitted time-generating function, supplies a functional relationship between line number and a present time of scan.

3.2 PREPARATION OF ADAS DATA

The block diagram for the preparation of the ADAS data is shown in Fig. 4. The program ADAS STRIP produces from ADAS data (the CCT is a Univac product) a set of records called ADAS. Each record contains a time reference as well as pitch, roll, and altitude information; the records are reported approximately 40 times per second. The timing of pitch, roll, and

INPUT DATA

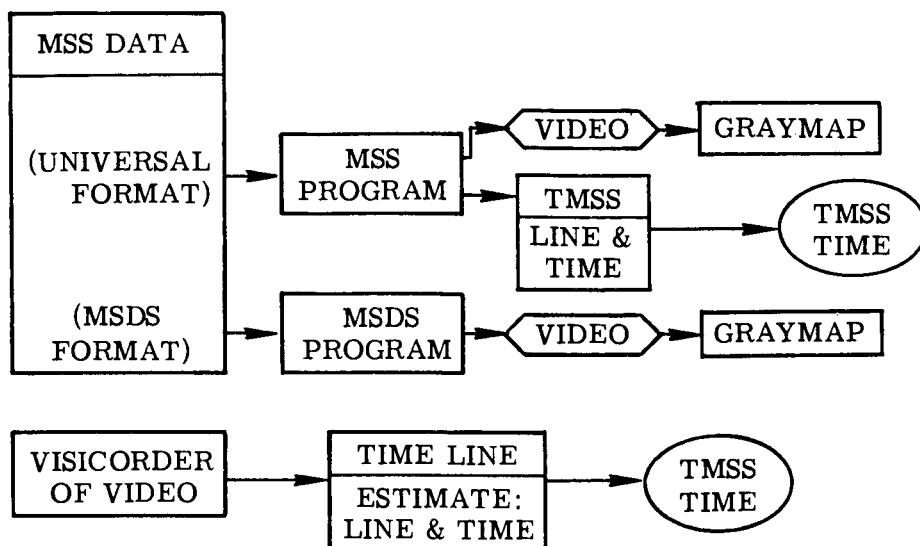


FIGURE 3. PREPARATION OF MSS DATA

INPUT DATA

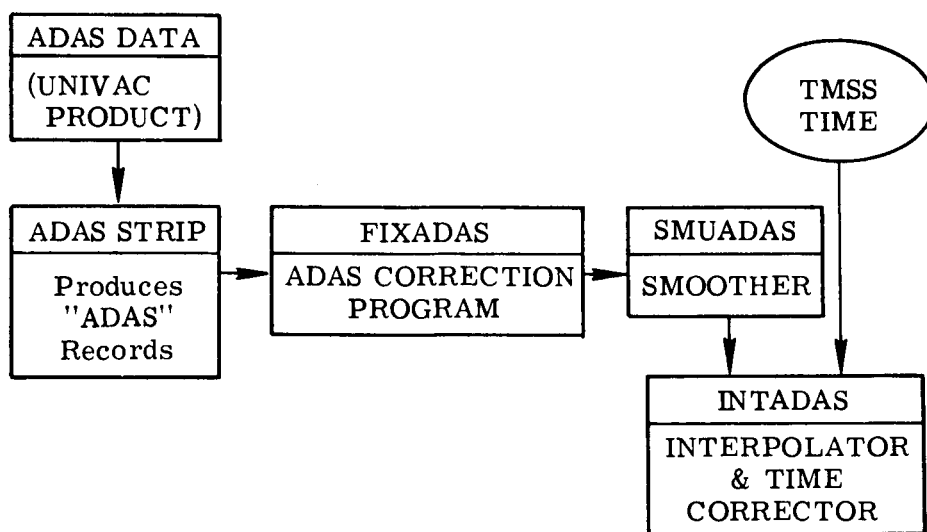


FIGURE 4. PREPARATION OF ADAS DATA

altitude has a small fixed offset from the reference time on each record. Pitch and roll are reported in degrees and quantized to a tenth of a degree. Altitude is reported in feet and quantized to the nearest 10 ft.

The program FIXADAS inspects the ADAS records for time gaps or irregularities and repairs the records by linear interpolation. ADAS data, reported at 40 times a second, are more than adequate to describe the highest frequency components of the aircraft motion. The data are marred by the least count, 0.1° , which is of the general size of a pixel, and thus may be expected to contribute round-off errors.

Smoothing these data appears worthwhile. The consequent loss in bandwidth is much less important than the reduction in noise and the smoothing over the stairsteps of the 0.1° least count. In a first reduction, a nine-point weighted moving average is used. The program to accomplish this smoothing is called "SMUADAS"; it is complicated by the amount of data to be processed, and by the necessity to look both ahead and behind in the data to form the moving average. The program reads in data in groups of 500 records, saves and moves data lines as necessary, and outputs smoothed data groups of 500. There is no reduction in the volume of the data in this processing, and SMUADAS will in the future be combined with the interpolating program which follows (see INTADAS in Fig. 4) to avoid the need for intermediate smoothed ADAS records. Program SMUADAS accepts the fixed ADAS records as input; it outputs smoothed ADAS records.

The next program, the ADAS Interpolation program, INTADAS, accepts the smoothed ADAS data and, in principle, gets scanline time (TMSS) from the video record. (Actually, along with the MSDS video tape, it synthesizes the time at which the scan line occurs.) Program INTADAS generates interpolated records containing scan time (TMSS), scanline number, and smoothed and interpolated roll, pitch and height of the aircraft. There is a reduction in the volume of interpolated data if the scan speed is less than 40 rps.

3.3 PREPARATION OF LTN-51 DATA

Preparation of the LTN-51 data is diagrammed in Fig. 5. After the program "LTN STRIP" is used on the LTN-51 data tape (Univac product), records called "LTN" are produced. Each record contains a double-precision reference time, the direction of the ground track, the heading of the aircraft, and the ground velocity of the aircraft. The angles are in degrees, reported to a least count of 0.1° . The velocity is reported in knots, with a least count of 1 kt. The time interval between records is about 1.8 sec, and a somewhat complicated situation exists in relating the timing to the data observations. An uncertainty arises because in the Litton navigation system the data are delivered to a buffer, which holds all the quantities and is updated at 1-sec intervals. The buffer is asynchronously read by the digital recording system, in fixed

INPUT DATA

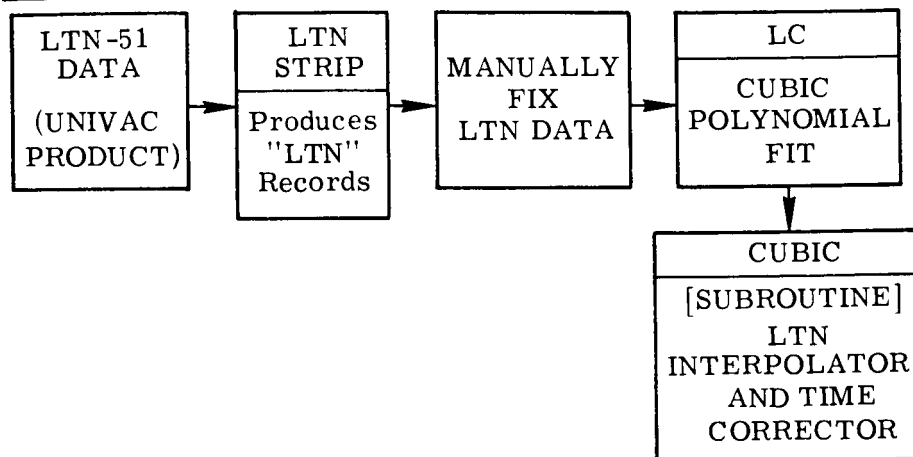


FIGURE 5. PREPARATION OF LTN-51 DATA

order, and reported. This results in a fixed delay which is known, plus a probable delay which is not known but expected to have a value of 1/2 sec.

The LTN data, recorded at 1.8 sec intervals, are sampled all too infrequently. To smooth these data would not be appropriate as the bandwidth is already inadequate. Instead, we must accept the least count errors and interpolate prudently between points. We have chosen to interpolate between four points with a general cubic equation, and to confine the interpolation to the (moving) region between the two central points. (As a special provision at the beginning and end of the data, interpolation goes all the way to the end points.)

To accept a wild point in this scheme of interpolation would be quite undesirable. It is thus necessary to inspect the data in the LTN record produced by the LTN STRIP program. The data are not numerous, hence this step has not been automated. It is presumed that an eyeball inspection is adequate, and easy. A wild point would give a contorted cubic function, from which even wilder points might be interpolated.

The main processing and correction programs need the heading data of the aircraft for scanner attitude estimation, and need it on a wide bandwidth basis (wider than now available). On the other hand, track direction and velocity are used only in estimating the present position of the aircraft and need not be wide-band. The desired output from these LTN data is really the present position of the aircraft, in some relative frame, not the track angle and velocity which are derivative quantities requiring integration.

Accordingly, the x and y positions are computed by integration of the track angle and velocity at 1.8-sec intervals in the LTN Coefficient program, (LC). (See Fig. 5.)

Within the LC program a cubic polynomial fit is made to all three input variables (i.e., heading, track angle and velocity). Thus the velocity, $V(t)$, is known by a function of the form

$$V = a + bt + ct^2 + dt^3$$

The track angle, TA is known by a similar function

$$TA = e + ft + gt^2 + ht^3$$

The positions are given by:

$$x_{n+1} = x_n + \int_0^{\Delta T} V \cos (TA) dt$$

$$y_{n+1} = y_n + \int_0^{\Delta T} V \sin (TA) dt$$

These integrals are unwieldy and have been replaced by:

$$x_{n+1} = x_n + (\bar{V} \cos(\overline{TA})) \Delta T$$

$$y_{n+1} = y_n + (\bar{V} \sin(\overline{TA})) \Delta T$$

where

$$\bar{V} = \frac{1}{\Delta T} \int_0^{\Delta T} V(t) dt$$

$$\overline{TA} = \frac{1}{\Delta T} \int_0^{\Delta T} TA(t) dt$$

The output of program LC is not a set of values for each pixel nor for each scan line but a set of cubic polynomial coefficients for an interpolating function. The three sets of coefficients associated with heading and position in x and y are inputs to the subroutine program called CUBIC.

In subroutine CUBIC, given scanline time (TMSS) as an input, the heading and position can be calculated as an output from the interpolating function while taking into account the time correction needed on heading and position data.

The position as calculated assumes north as signaled from the navigational equipment, and takes as zero position the arbitrary position of the aircraft at the start of the LTN data. No allowance has been made for the convergence of the lines of longitude as the effect is of marginal importance in the short intervals of data taking; such allowance could, however, be readily included in the program.

3.4 PREPARATION OF GROUND TRUTH DATA

The "ground truth" record must be prepared from the MSS data by comparing graymaps with photographs and maps. The record then consists of the scanline and point numbers corresponding to landmarks in the graymap, together with the locations of the same landmarks in map grid coordinates. The ground truth record ideally should contain a good statistical sample of points, well distributed within the scene.

The MSS data for this study were collected over the Casa Grande test range in Arizona which has surveyed landmarks (concrete squares with Maltese crosses) forming a 16 x 16 mile grid network with 1-mile spacings. In preparing the ground truth record, we now find that it may be desirable to supplement the record of the surveyed landmarks with as many other landmarks as can be recognized and located in both the graymaps and maps (or photographs). In practice, of course, unsurveyed landmarks are likely to be the only type available for ground truth.

3.5 PRELIMINARY GEOMETRIC CORRECTION AND FITTING

The preliminary geometric correction and data fitting functions, represented in the flow diagram (Fig. 2) by two blocks so identified, are handled by one program—program MAP. This program is a central one in which the four types of data converge and a fit to the ground truth data is obtained.

The program is given the scanline and pixel numbers which coincide with each ground truth point. It reads in the smoothed and interpolated ADAS data in blocks, and selects the data line matching the scanline number. (The program cannot go backward in the ADAS data, thus the ground truth data must be ordered in terms of scanline number.) The program also enters the LTN coefficient data, selects the appropriate set of coefficients, and interpolates to find the LTN values. Somewhat intricate timing matters are considered in bringing the several records into accord.

For the first time a full set of position and attitude data is now assembled pertaining to a scan line. Calculating the ground point to which each pixel refers at last becomes possible. This is done for each of the ground truth pixels, and a set of provisional locations is arrived at for these points. The convention is adopted that the aircraft is not changing attitude during the scan; rather, it is considered held in the attitude it had at the mid-point of the scan. The aircraft is considered to be moving forward during the scan, however, in accordance with its velocity vector.

This step in the program constitutes the preliminary geometric correction made in the ground location of a pixel based on the input aircraft information. The x and y positional equations relative to the nadir are

$$x = H[\tan(\phi + \lambda) \sin \psi \sec \theta + \tan \theta \cos \psi] + Vt \cos \beta$$

$$y = H[-\tan(\phi + \lambda) \cos \psi \sec \theta + \tan \theta \sin \psi] + Vt \sin \beta$$

where

ϕ = roll angle

λ = scan look angle

ψ = heading angle

θ = pitch angle

β = track angle

H = height

V = ground velocity

t = time reference from midscan

The position of the nadir is obtained from the LTN record. (For axis system conventions, see Appendix.)

The preliminary corrected ground points are now available for comparison with the ground truth positions. The ground truth positions are in the (metric) UTM grid system, and thus have a coordinate zero determined by the cartographer. The calculated positions are based on a zero at the start of the LTN data. There is thus a large and arbitrary disagreement in the location of the origin. The ground truth values are in terms of a local "grid north," which may depart somewhat from true north; but the north of the LTN navigator may be technically in error, thus a disagreement between the reference directions may be expected as well as a disagreement in the zero positions.

Other errors may be expected in the data. The LTN equipment must be boresighted with the scanner. A procedure exists for aligning the LTN with the airframe, which involves jacking the aircraft into level attitude, etc. We are told that LTN units are frequently exchanged for maintenance. Thus it seems beyond the bounds of the probable that the LTN units are consistently in alignment with the airframe.

The airframe alignment marks are in fact not a guarantee of the attitude of the aircraft in flight. Rather, the attitude in flight depends on the loading and weight, and on the airspeed; it also depends on the trim settings. The weight changes as fuel is consumed.

The installation drawings for the scanner make no special mention of a precision required in alignment, nor have we determined that the device was aligned by survey. Survey would be difficult as the scanner sits at the bottom of a well. It is only reasonable to assume that the scanner is aligned with the accuracy of a sheet-metal worker's eye. In any case it sits on shock mounts, so that the final alignment depends on the weight distribution and on gravity. This device, which must have an appreciable gyroscopic moment, is used in an accelerated environment. Those shock mounts are unfortunate.

The LTN unit is well aft in the aircraft, while the scanner is forward. Aircraft are elastic. A survey made while the aircraft sits jacked level on the ramp would really not pertain to an aircraft in flight. In flight the weight is carried on the wings instead of on jacks and wheels, engine thrust is powerful, and there is a strong downward force acting on the stabilizer. Obviously, the static force system with the aircraft parked on the ramp is quite different. We understand that in surveying the aircraft, the fin may be used as a center-line reference. Characteristically, the fin is not on center; propeller torque is balanced by offsetting the fin.

Everyone has observed the dynamic deformation of aircraft fuselages and wings in flight. The torquing of a conventional fuselage is easily (or uneasily?) visible to the eye, and amounts to several degrees. We once made observations on a cargo plane with a large rear door and high tail, and a twist of several degrees was observed. The consequences of such flexure in flight will be unfortunate in the present work, yet a remedy is available: If the LTN equipment were installed on the scanner, there would be no trouble.

But even if that were done, calibration problems would still be present. In particular, the scanner time reference, which indicates where the scanner is looking at a particular time in the scan, seems uncertain. To determine this in ground testing appears most difficult. A remaining uncertainty is whether the scanner is indeed looking at right angles to the spin axis. If it is not, the consequences are not pleasant. In this particular matter we place our faith in the precision of the optical instrument maker.

Because this is really a surveying job, we are interested in picayune details of precise alignment. We must plan to overcome the shortcomings of the data as best we can. It appears that we have a rather rubber-legged transit about which little can be done, yet we hope to be able to offset the systematic errors of alignment of the system. In any case, a certain amount of ground truth data must be furnished. By supplying enough ground truth points to permit a statistical fit, we believe it possible to remove all the calibration errors from the system. The method is not quite that used in cartographic fitting, but special.

Of the errors in the system, two are particularly troublesome. One is the possible fault when the scanner look direction is not normal to the scanner axis; this results in the "plane" of scan becoming a cone. Taking the x direction as that of heading, the result is a hyperbolic intersection of the scan cone with a flat earth. An error is introduced in x that is proportional to the square of y—i.e., a "quadratic error". As noted above, we take it on faith that the scanner does not have this fault, yet it would be possible to examine the data with a special program to see if this is true.

The second troublesome error is an error in timing, leading to an error in knowing the phase angle of the scanner rotation. This also results in another quadratic error, an error in y proportional to y squared.* We do presume this error to be present in the data, if only because of our inability to discover the precise convention used in the timing, and to know how boresighting is accomplished.

The y mentioned above is not the y coordinate of the located point. It is instead measured in a system rotated to the aircraft heading, and is a local y_{offset} from the present position of the aircraft. If the coordinates are computed in a north-oriented system with a fixed zero, this y_{offset} is not an available quantity. A standard cartographic fit cannot remove this error. Hence, we are forced to use a new coordinate system, aligned with the aircraft axis, and to carry y_{offset} as an extra quantity. The computing effort to do so is not great during fitting as only the few ground truth points are involved, though it would become a burden in processing voluminous scanner data. Our tactic is thus to go through a preliminary statistical fit specifically to learn the error in the scanner angle. We then correct the scanner angle, compute a new set of ground coordinates, and do a standard cartographic 2×2 fit.

*The genesis of this quadratic term will be explained later.

The process is as follows:

(1) The estimated x and y values of each ground truth point are calculated in a rotated coordinate system in which the x-axis lies along the mean heading for the run. The mean value of x is subtracted from the x values, and the mean y is subtracted from the y values. A y_{offset} is also calculated for each point.

(2) The mean values of the true x and true y data are subtracted and are similarly rotated. (It is essentially an accident of programming that the sequence is different, translation before rotation, but the intent was to preserve precision. The grid values may be very large numbers, and it is well to subtract the mean before any multiplications. The resulting means of the estimated and true values are hence not in the same coordinate system!)

(3) Error correction is accomplished by the following two formulas, with the K values to be determined by a statistical fit for minimum total squared error:

$$x_{\text{error}} = K_{11}x + K_{12}y + K_{13}(y_{\text{offset}})^2 + K_{14} - x_{\text{true}} \quad (1)$$

$$y_{\text{error}} = K_{21}x + K_{22}y + K_{23}(y_{\text{offset}})^2 + K_{24} - y_{\text{true}} \quad (2)$$

The error quantities represent the difference in position between the calculated x,y and the true x,y for each ground truth point. To solve for the K's, square both equations, and sum over all the ground truth points. Take the partial derivatives of each resulting equation with respect to the K's contained in it. Set each partial derivative equal to zero, representing minimum over-all error. The results are two independent sets of four simultaneous equations, which were solved by matrix methods. The matrix of coefficients consists of sums of products of the geometric variables, and is common to both sets. The vector of constants consists of sums of products containing the true variables.

The matrix was inverted via the routine MINV of the IBM scientific subroutines package, and the inverse multiplied by the vectors of constants. This works successfully but we remain somewhat uneasy since the matrix is quite "ill-conditioned" in that the elements on the principal diagonal are disparate in size, varying by a factor of 10^{10} . There may be loss of precision with some data, and double precision may be called for at this step.

From the fitting matrix we are interested only in the coefficient K_{23} , which when multiplied by the aircraft height gives the scanner angle error in radians for correcting the data. The relationship is explained as follows: In the rotated (with respect to mean heading of aircraft flight) coordinate system, y_{offset} is the tangent of indicated scanner angle (which has been corrected for the roll angle of the aircraft) multiplied by a quantity which is nearly the height of the aircraft. The true y_{offset} is similarly found, but from the true scanner angle.

The error between these two y offsets is thus proportional to the height, but also proportional to the difference in the tangents of the indicated and true scanner angles. Let us express in a Taylor series, the tangent of true scanner angle in terms of the indicated scanner angle and a small angular error. The Taylor series includes tangents as well as secants which can readily be expressed as tangents. The Taylor series for the error can thus be expressed in terms of y_{offset} since the y_{offset} is proportional to the tangent of the indicated angle. Using this expression and neglecting higher order terms, we can compare it with the expression for the y_{error} (Eq. 2) in order to identify the coefficient of $(y_{\text{offset}})^2$, K_{23} , as the scanner angle error divided by the height.



It may be erroneous to attribute causality, yet if we anticipate scanner angle error, and find error in the data of the sort that modification in scanner angle can eliminate, the change in scanner angle should provide practical data correction—which is the goal.

It should be possible to attribute causes to the other elements of the 2×4 fitting matrix, but the matter is complicated by the particular gimbal sequence, and in fact some of the errors can have more than one cause. For instance, removal of the scanner angle error alters not only the linearity of the image, but also its position and orthogonality. (The scanner angle occupies a position in the gimbal sequence that permits ready interpretation of K_{23} .) To get around these complexities we adopted a sure way, which offers a check on the correction of the scanner angle as well. We are in search of a practical 2×2 matrix fit. We corrected the scanner angle, then made a new x and y estimate and a 2×2 fit. We examined the residual errors in this second fitting and found them the same as in the earlier 2×4 fit. The program was tested by inserting a scanner angle error of 0.1 radian, or about 5° . This badly spoiled the initial fit, but the correction process smoothly removed it. Limited testing with other types of error has also been conducted. The fitting procedure makes all the necessary corrections. One type of error has not yet been tried—namely, an unknown altitude error.


This is essentially the end of the central program "MAP." What has been achieved by the program is a means for calculation and correction of a ground location, given a scanline number and a pixel number on the scan line.

Program MAP prints, for inspection, the ground truth points at several stages of the calculation, and also gives the rms error at each stage. Useful outputs are: a 2×2 matrix for correction, the necessary translation of the data, and the scanner angle correction value.


The types of errors which a 2×2 fit removes are these:

□ This is a true square on the ground. This square may be seen in the data as rotated: ; the 2×2 fit rotates it back. It may also be seen in the data as out of scale: ; the 2×2 fit pulls it into shape and into scale.

It may be seen in the data as out of square:

 The 2x2 fit straightens it out.

All the above errors can happen at once:

 The 2x2 fit fixes them all at the same time.

One may identify causes of some of these errors. The rotated square, for example, comes from error in LTN north or deviation of grid north from true north. The parallelogram distortion can stem from an unrecognized yaw or drift, or a yaw boresight-error in the scanner. The scale error can come from height and velocity errors, or from unrecognized pitch errors.

However, if the 2x2 fit is applied to the data—and this is the simplest method—it is not necessary to identify causes. (Attribution of causes necessitates establishing a sequence for removal of errors which are not independent of the sequence. It is tidier not to attempt to attribute causes.)

3.6 FINAL DATA PROCESSING

Thus far we have discussed only a part of the problem, namely the assignment of a true position to each picture element. The remaining portion of the problem is to assign an area to each picture element; then, by attaching a recognition algorithm, the acreage of a given crop can be measured. In order to be able to say where—that is, in what parcel—that acreage lies, we must be able to describe and use a set of boundaries. We must also find more efficient means for assigning locations to the picture elements.

To handle the problem of boundaries efficiently, we must deal not in points, as was done in program MAP, but in scan lines. We may reasonably ask whether a scan line enters a field, and if so where. We then identify the point numbers associated with the line intersection with field boundaries. We may then "enable" the recognition algorithm and the area algorithm and find that area within the field devoted to a particular crop (or crops).

In the several programs the scan lines as projected onto the ground have been described by their equations. These lines, if the scanner "looks" normal to its axis, are straight except for the correction due to aircraft velocity. The amount of nonlinearity introduced by aircraft velocity is negligible. The velocity merely increases the slope angle of the lines by a small amount, and makes a slight change in their intercept.

A program addressed to the boundary problem has been written. This program, called "EDGES," accepts four points sequentially bounding a convex quadrilateral. It tests to see that the quadrilateral is indeed convex, computes the lines bounding the quadrilateral, and tests scan lines to see if they enter. If they do, the program produces the pixel point numbers of entry and exit. An algorithm not yet added will find the true area of the quadrilateral. (Program EDGES, though written, has not been debugged; it should prove useful in future development.)

A second reason for expressing the data in terms of scanline segments lies in the algorithm chosen to calculate the area to be assigned to each pixel. The actual area the scanner instantaneously sees on the ground is the area upon which any recognition is based. This is not the area which must be assigned to a pixel, however. Thus the actual resolution elements will overlap in varying degree, particularly toward the sides of the picture. But the increments of area which must be assigned these elements are instead chosen to be such that their total, within a field boundary, equals the geographic area within that boundary.

The area algorithm considers three adjacent scan lines, as shown in Figure 6. The n -th scan line has large dots added, representing the location of some adjacent pixels along that line. The dashed lines divide equally the space between adjacent scan lines. In implementing this algorithm it is convenient to use the solid lines paralleling the x -axis to demarcate two sides of each pixel. The quadrilateral formed by these dashed and solid lines (see shaded portion of Fig. 6) represents an increment of area to be assigned one pixel. The pixel area is the x height of the zone times the y width. Such increments when summed, will equal the total area scanned.

Each scan line is represented by $x = my + b$ in the coordinate system. It is evident that the x height of the shaded zone is half the x distance between the $n + 1$ scan line and $n - 1$ scan line, or

$$\frac{x_{n+1} - x_{n-1}}{2} = \left(\frac{m_{n+1} - m_{n-1}}{2} \right) y + \frac{b_{n+1} - b_{n-1}}{2}$$

If the y value of the q -th pixel is represented as y_q , the area to be assigned a pixel is:

$$\Delta A_n = \int_{y_{q-1/2}}^{y_{q+1/2}} \frac{x_{n+1} - x_{n-1}}{2} dy$$

where $\frac{x_{n+1} - x_{n-1}}{2}$ is function of y_q .

To assign an area to each pixel and to then sum these areas, however, is not an efficient process. If we are in the middle of a wheat field, the most likely case is that the adjacent element also contains wheat. We will find continuous strips of wheat recognition, from say point a to point b . In such case the increment of area is found more efficiently by:

$$\Delta A_{ab} = \int_{y_{a-1/2}}^{y_{b+1/2}} \frac{x_{n+1} - x_{n-1}}{2} dy$$

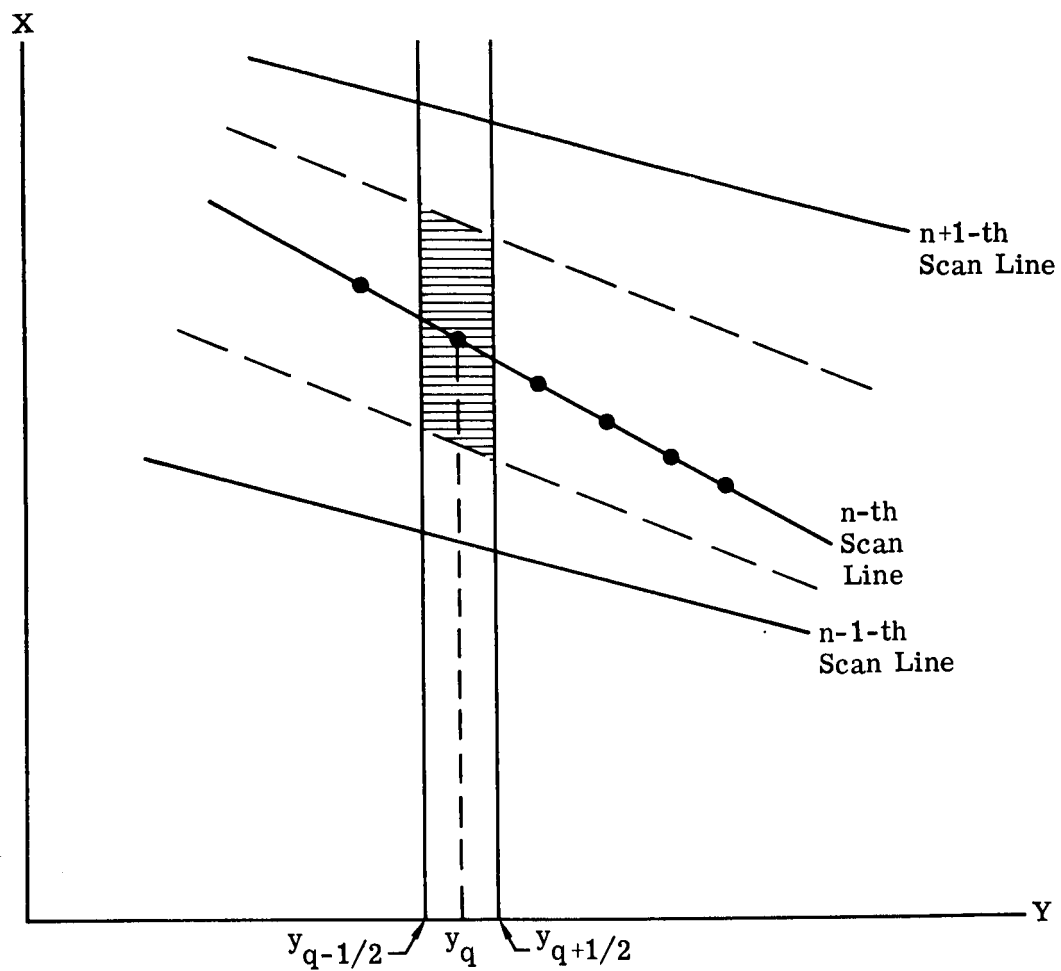


FIGURE 6. ALLOCATION OF PIXEL AREA USING AREA ALGORITHM

This in effect takes long strips as the increments of area. To use this shortcut we need an "abstract" of the video recognition results reading, in part, something like this:

"wheat" from pixel a to b

" ? " from pixel (b+1) to c

"wheat" from pixel (c+1) to d

.
.

and so on.

Now if every adjacent pixel shows a different recognition, the area calculation devolves into adding the increments one at a time, so nothing is lost. Efficiency is gained if wheat appears beside wheat, for example.

The area algorithm discussed above has been embodied in a final data processing program. This program, described in the next section, has been used and tested on simulated data.

By using simulated data in the system, we can demonstrate that the essential programs of the system do indeed work. This is expedient since processing of existing live video has not progressed far enough to permit system testing. In any case, live video tests could prove to be a difficult means of confirming the area algorithm unless boundaries of a known ground area were to be manually selected and specified in graymap point and line numbers. Moreover, the real data contain random uncertainties in timing which could defeat confirmatory testing of the programs.

Because it was foreseen that real data might not be on hand soon enough to use in the overall program development, four simulators were programmed to generate synthetic records of video, ancillary equipment (ADAS and LTN-51), and ground truth data. In addition to these four simulation programs, another program was produced which is essentially a prototype of the Data Processor block for real data; this program was run to demonstrate that the entire structure does indeed function.

For demonstration purposes, an aircraft flight path and attitude regime was assumed that could be described by functions of time. The aircraft, flying at constant velocity, is placed in a very slow turn to provide an accelerated path; made to roll back and forth sinusoidally; given a steady drift; and made to yaw one way and the other in the horizontal plane. It was also made to pitch up and down cyclically. Altitude (height above terrain) as well as speed were constant.

Since the attitude and position of the aircraft are described by exact functions, any of the simulators can calculate exact data from time alone.

4.1 ADAS SIMULATOR

The ADAS simulator, named ADAGEN, calls on pitch, roll, and height functions and generates simulated ADAS data, quantized to 0.1° increments in the angular quantities and with a reported time offset by the proper amount. Provision exists for inserting fixed errors. (Random errors could also be included but have not been.)

4.2 LTN-51 SIMULATOR

Our second simulator, called LTNGEN, generates synthetic LTN-51 data. Using heading and velocity vector functions, the generator produces data quantized in 0.1° intervals and with the systematic time offsets characteristic of LTN data. The random time uncertainty present in the real data is not included, however, as our interest in the correctness of downstream programs is better served by the absence of such uncertainties. There is provision for inserting systematic errors.

4.3 GROUND TRUTH DATA SIMULATOR*

The Ground Truth Data Simulator is more complicated. Because it simulates what an actual scanner would see, it does not include errors. (Such errors occur in the reported altitude, pitch, roll, etc. of the aircraft, not in the actual values of such parameters.)

Ground truth points were generated as follows. First, six conceptual lines were layed out on the ground, running north and south (x direction). These six marker lines are spaced at equal intervals from $y = -2000$ to $y = +2000$ meters. The simulated aircraft path, though curved, was generally northward, so the six lines remain more or less within the field of view of the scanner. (The flight, as simulated, lasts 50 seconds and is 1001 scan lines long.) Ground truth points were taken at the intersection of the first scan line with each of these six ground lines, and thereafter at intervals of 50 scan lines to the final line number 1001; this resulted in 126 ground truth points. The coordinates of the intersections were found and recorded. The scan-line number was recorded, and the nearest integer pixel number was recorded.** Note that this falsifies the data slightly in that the ground truth values are always exactly on a scan line, not just on the nearest line. (This does show up in the processed data where errors in x are less than those in y.)

4.4 VIDEO SIMULATOR

The Video Simulator is named VIDGEN. The problem was to lay out a pattern on the "ground," fly the simulated scanner past it on the wobbly path described by the sinusoidal functions, and produce the video the scanner would have seen. As in the ground truth simulation (and for the same reasons), there are no errors in the aircraft path or the video. Instead, the errors are all in the reported altitude, pitch, roll, etc., of the aircraft.

The pattern on the ground represents simulated recognition. The machine is fed unambiguous recognition of 'yes' within the limits of the pattern and 'no' outside the pattern. This solves the boundary dilemma. Since the true area is known, the object is to find the area by means of the area algorithm.

The first task was to choose a suitable ground pattern. Constraints were that the pattern must fit on a printer plot and yet must be large enough to reveal errors in the process. To make the choice more difficult, a well dissected outline is desirable, not a simple rectangle. The pattern should be one that is readily recognizable, yet feasible to program.

*The ground truth simulator uses geometric equations for scan lines rotated in a somewhat approximate manner to represent the effect of aircraft velocity, whereas program MAP deals with the geometric equations for points which are advanced in an exact manner to account for aircraft motion.

**No angular limit was placed on the scanner, hence a few fictitious pixel numbers under 1 or over 700 may have been generated.

Something in block letters would be ideal. It is regrettable that the initials of our Sponsor as well as those of our own organization contain letters having diagonal strokes and hence had to be rejected. The pattern chosen was the word "TEST."

The letters were assembled from modular squares, each measuring 120 meters on a side. The letters were made 3 modules wide by 5 high, and were separated by a space 1 module wide. The bar of the "E" was made equal in length to the top and bottom strokes. Each module has an area of 1.44 hectares. (A hectare is about 2-1/2 acres, and exactly 10^4 square meters.) The block letter "tees" thus have an area just over 10 hectares (10.08, to be exact), or about 25 acres apiece. The letters "E" and "S" each have an area of about 40 acres, or 15.84 hectares. The total area is 51.84 hectares, or about 130 acres, which is somewhat under the size of the usual 180-acre homestead. This pattern affords a difficult test of the area algorithm. No farmer would lay out a 40-acre crop pattern in the form of the letter E, for example.

When corrected, the resulting pattern should just fit on a printer-plot page, in maximum expansion and with 6×10 meter resolution cells, which gives an equal scaling of 60 meters per inch in both x and y.

The pattern of the work "TEST" was enclosed in a rectangle 600 meters wide by 1800 meters long, and this was divided into 5 squares in width by 15 in length; then a border of blank squares was added all around the edge, making a 7×17 matrix. The inner portion of the matrix was given values of one and zero, with "ones" forming the word. The upper left block of the T is in element 2.2 of the matrix. Two guidelines were placed on the ground, one at $y = -600$ meters and the other along the x-axis. These demarcate the top and the bottom of the lettering.

The program is instructed regarding the range of scan lines to generate, and consecutive scan lines are generated over this range by the method used in the ground truth generator. The pattern is approximately centered in the 1000-line run at $y = 3050$ meters, and lies between $x = 2150$ and $x = 3950$ meters. Thus, if the values of the intersections of the scan lines and the guide lines are substantially less than 2150 meters, blank video is generated, and the same is true if the intersections are substantially greater than 3950 meters. Between those limits, however, picture video is produced. (The switching on and off of the video generator at a geographic boundary—not at a line number and pixel number—anticipates the future objective of being able to observe a described geographic area.)

The video consists of close-packed binary bits, "one" for points lying inside the letters and "zero" for points outside. This packing is eight times more dense than in the real video, which is permissible as the video represents recognition rather than a gray level. One computer word represents 32 pixels and 22 words represent a scan line of 700 pixels. Each scan line is given a line number; the scan time for each line should also be given but is not, as these times are missing in our real video. This computer word format is compatible with the IBM 360, which has 32 bits per word.

The program first locates the pixel or "point" number corresponding to $y = -600$ meters and also that point number corresponding to $y = 0$ meters, and then slightly increases the number of points considered. From pixel 1 to the point of beginning, the "packer" now packs zeroes. At the point of beginning it computes the pertinent matrix element by dividing x and y by 120 and converting to integer. If the point lies outside the range of the matrix, it is placed in the border of the matrix. The matrix is then consulted, and it returns a one or a zero depending on whether or not the point is within a letter. This zero or one is packed into the video. When the far edge is reached, the packer fills the remainder of the 22-word video line with zeroes.

This pattern would be too large to fit on the page printer, requiring some six side-by-side printer pages. It can be examined readily, however, by listing the file containing it in hexadecimal, where one printed character represents four bits of the packed video according to the code:

0000	= 0
0001	= 1
0011	= 3
0111	= 7
1111	= F
1110	= E
1100	= C
1000	= 8

To also compress the video by a factor of four in the direction of travel, every fourth line of the file may be copied. Figure 7 shows such a copy which, to enhance the contrast, has been edited by replacing all the zeroes with spaces. Because the scanner rotation direction is such that the image would be in mirror writing, this has been avoided by copying the file in reverse, last line first.

The resulting video picture is violently distorted. The aircraft was yawing, pitching, rolling, and turning, though in a known manner. The only real error introduced is that resulting from the quantization into pixels.

4.5 PRELIMINARY GEOMETRIC CORRECTION AND FITTING

The program 'MAP' was used on the simulated data. Two runs were made with simulated data. In the first run, the only errors introduced were in round-off quantization of the simulated ADAS and LTN-51 data. In the second run, the following systematic errors were also introduced:

Δ Velocity	= 10 knots
Δ Pitch	= $+2.5^\circ$
Δ Roll	= $+5.0^\circ$
Δ Heading	= $+5.0^\circ$
Δ Track Angle	= $+5.0^\circ$

Details of this step of the processing are covered in Section 3.5. The 2×2 matrix values of the standard cartographic fit were separately calculated and used in the final data processing for each run.

A set of simulated data was initially used to test the program MAP by inserting a scanner angle error of 0.1 radian, or about 5° . The initial fit made by the preliminary correction gives poor results, as would be expected. The interpretation of the K_{23} coefficient in the 2×4 matrix fit corrected the error. Time has not permitted exercising the fitting program with all types of errors. Because the simulator programs were written so that either constant or random errors can be included, the simulator will continue to be a useful tool in future work.

4.6 FINAL DATA PROCESSING

This last program of the main series is a simulated "using" program called OUTLINE. In essence it is not a simulator, but a working program that uses the simulated video. Written in a modular form, it has subprograms which can be re-arranged to perform other "user program" tasks. The program computes the area of the word "TEST," and as a by-product, outputs a corrected image showing the outlines of the word "TEST."

The program has a special limitation which need not have been embodied; this has to do with scanline equations in the corrected (or grid) coordinate system. This will cause the program to fail on easterly or westerly headings. The program goes to some pains to get into this dilemma, rotating the 2×2 correction matrix into the true coordinate system, etc. We might better have worked in the coordinate system rotated by the mean heading, where the problem cannot arise. However, because the simulated video was taken on a northerly heading, there was no problem.

The program corrects the scanline equations using the 2×2 matrix. The scan line has one singular point on it: the location of that point at which, in the scanner plane, the scanner beam is perpendicular to the intersection with the ground. This singular point is given two descriptions, one in terms of the scanner resolution element number at which it is met, the other in terms of the x,y coordinates on the ground at which it is met. The resolution element number is treated as a real number, not as an integer, and includes correction for roll and for the scanner angle error. The x,y coordinates of the center are corrected by the 2×2 matrix. The slope, intercept, central point number, and central point location afford a compact description of an entire scan line.

Preparation of still another intermediate record would be logical at this point, for practical use. Such a single record, plus the appropriate algorithm, could be used by a recognition program. That course was not adopted, however, as the system seems to have more than enough intermediate records already. The program instead calculates the scan lines as required and handles the LTN coefficient data, the interpolated ADAS data, the video data, and certain nec-

essary control data all at once. It writes the computed area and possible error messages on one output device, and the corrected map on another.

A subroutine prepares an abstract of the video, generating the numbers of "recognition blocks" of "ones" found on a scan line, and the point numbers at which they begin and end.

Another subroutine was written to read-in the interpolated ADAS data in blocks and search out the desired line number. (This subroutine could have been used in MAP.)

A variant of a subroutine previously used in MAP interpolates to find the LTN data.

A subroutine LINES then calculates the corrected scan line on the basis of the interpolated ADAS and LTN-51 data. The data are not looked up nor is the line calculated on scan lines for which no recognition blocks were found. From the video abstract the corrected ground positions of the block edges are calculated. (For reasons of efficiency the program does not deal in points internal to such a block.) These edge points are recorded for subsequent plotting via a printer plot routine.

It will be remembered that the area algorithm requires the coefficients of the prior line and the subsequent line. The subprogram "LINES" thus normally stays one line ahead of the video. There is an exception to this when one or more blank lines occur; here the subprogram falls behind and has to catch up again when more video is encountered. Values for previous lines are stored in a bottomless push-down stack prior to the calculation of a new line.

The area algorithm previously discussed is now used, and the area is recorded as a sub-sum every time a gap is found between simulated letters. This feature works only on the initial "T," because the aircraft attitude is such that there are no empty scan lines between the subsequent three letters (see hexadecimal printout of simulated video, Fig. 7).

After OUTLINE has been executed, a very small program produces the printer plot from a list of previously computed and filed coordinates.

4.7 RESULTS

The corrected image of the simulated video shown in Fig. 7 is presented in Fig. 8. The letter outlines are no longer distorted and the areas computed are correct. This example demonstrates the ability of the SICLOPS System to place resolution cells in their correct location and to compute ground areas using the area algorithm.

Results were as follows:

(1) In the first run, there were round-off and quantization errors only in the simulated ADAS and LTN-51 data (see Table 1).

- a. The calculated area of the initial "T" was 10.05 hectares. The true area was 10.08 hectares.

b. The calculated area of the entire word "TEST" was 51.72 hectares. The true area was 51.84.

(2) In the second run, there were systematic errors in addition to round-off and quantization errors (see Table 2).

a. The calculated area of initial "T" was 10.10 hectares.

b. The calculated area of entire word "TEST" was 51.99 hectares.

(3) In both runs the pattern of the word "TEST" is without distortion, and the boundary has only a few displacements; these were by no more than one plotting cell. The cells are 6×10 meters. Because the outlines for both runs were so similar, only the outline for the first run is shown in Fig. 8.

(4) The position of the word is correct in both x and y.

(5) Location accuracy is indicated by the rms errors as calculated in the map fitting program. The radial rms error is essentially 4 meters in Run 1, and 6 meters in Run 2. The errors in the first run may be attributed primarily to round-off errors in the data and to quantization of the video into pixels. The added component of error in the second run may represent primarily accumulated computational errors.

TABLE 1. AREA MEASUREMENT OF "TEST
PATTERN ON RUN 1. (Includes round-off and
quantized errors.)

<u>Subject</u>	<u>Area (Hectares)</u>		<u>% Error</u>
	<u>Actual</u>	<u>Computed</u>	
Word "TEST"	51.84	51.73	-0.21%
Initial "T"	10.08	10.05	-0.30%

TABLE 2. AREA MEASUREMENT OF TEST
PATTERN ON RUN 2. (Includes systematic
errors in addition to round-off and quantiza-
tion errors.)

		<u>Area (Hectares)</u>		<u>% Error</u>
		<u>Actual</u>	<u>Computed</u>	
SUBJECT	Word "TEST"	51.84	51.99	+0.3
	Initial "T"	10.08	10.10	+0.2

NOTE: The systematic errors were:
 Δ Velocity = 10 knots, Δ Pitch
= 2.5° , Δ Roll = 5° , Δ Heading
= 5° , and Δ Track Angle = 5° .

COMMENTS ON APPLYING THE SICLOPS SYSTEM TO REAL DATA

Because of limitations in time and funds it was not possible to prepare the ground truth record necessary to exercise the programs with the real data. The results that might have been obtained, however, had the real data been processed, are interesting to consider:

(a) In the real LTN-51 data there is a random uncertainty in timing. This is of consequence primarily in the heading information, and cannot fail to impair the accuracy of location. The random displacement of the values in time will not be treated well in the cubic interpolator, one may be sure. The data may have been damaged technically beyond any real repair. On the other hand, aircraft are conservative in their heading angle so the heading change may not be large.

(b) The heading, reported in the LTN-51 record at approximately 1.8 sec intervals, is not frequent enough to reconstitute the perturbations (in yaw) about the heading.

(c) The real ground truth data cannot be as precise as that used in the simulation. In the simulation there was no quantization error in x (as mentioned previously in Section 4.3) whereas in real data there will be errors in reading coordinates and in recognizing the landmarks. Glaring errors may be spotted by obtaining a printout comparing the computed positional coordinates with the true coordinates of selected landmarks. From this, genuine wild points will be recognizable and can be corrected or eliminated. Still, however, a substantially larger uncertainty than was simulated will remain.

(d) The effect of static flexure of the aircraft will be eliminated, but the effects of dynamic flexure cannot be, and may seriously influence the map. The result will be evidenced by residual local stretchings and distortions of the image.

(e) Except for rms errors found in the map fit, an objective measure of precision would be singularly difficult to obtain from the real data. However, a corrected graymap could be produced via the printer plot routine; from this printout, a good idea of the remaining errors might be gained.

CONCLUSIONS AND RECOMMENDATIONS

In the initial study of the scanner cartography problem, the quantitative effects of aircraft instabilities and scanner geometry on scanner data were identified, correction techniques utilizing information from ancillary sensors and ground truth data were specified, and an area algorithm was developed. The results of this study led to the development of a system of computer programs referred to as the Scanner Imagery Correction and Line-by-line Orientation ProgramS (SICLOPS) System. The system of computer programs as developed under the present contract has been tested with simulated data. Results indicate that the SICLOPS System works. The demonstration of corrected imagery supported by the area measurements results strongly verifies the feasibility of this approach for post-flight processing of scanner data. The SICLOPS System still needs exercising with real data.

Those system programs presently devoted to data preparation are specially designed to accept certain data provided by NASA's NC130B airborne system. These data come from the Multispectral Scanner, the Litton LTN-51 Navigation System, and the radar altimeter. The main programs of the SICLOPS System, however, apply to any imaging airborne scanning system and can be utilized as long as the necessary attitude and position information is furnished in time-synchronization with the scanner data.

The first real data assembled for processing by the SICLOPS System does not have the preferred MSS data format. The system was designed in the expectation that the data would be in the Universal format which includes a reference time for each scan. Instead, the real data actually to be processed (MSDS format) has no such reference time; its estimate by an auxiliary operation will thus be necessary. Results obtained on the same flight line for these two different MSS inputs to the system should be compared.

Combination of the SICLOPS System with the EDGES program and a recognition algorithm is an obvious next step. To estimate the area of a crop in a particular field by automatic means would then be within reach.

Elevation variations within the scanned scene are not handled in the present SICLOPS System; the severity of the problem should be studied. Possible methods of providing elevation information to the system should also be studied. One approach would be to use topographic information as it is known with respect to a geographic grid system. Another approach would be to consider modifying the scanner system to include a pulse laser to provide the slant range in time-synchronization with the scanner data.

Practical means for supplying the ground truth (control) data should be considered. To work from visicorder pictures, etc., rather than from digital graymaps, may prove possible.

Finally, experience could be gained toward the ideal of calibrating the static errors out of the data-taking equipment. If that were possible, the requirement for ground truth would shrink to the minimum necessary to adjust the coordinate zero, adjust scale, and rotate to grid north.

The SICLOPS System as developed thus far is not a final product. Its programs were developed sequentially as a programming convenience and thus provided an economy in computer usage. In consequence, some of the later programs have good features not found in the earlier ones. The present programs merit further clean-up to improve speed and ease of processing. Telescoping of several programs into one is possible, and the ease of operating the programs could be improved. Further development should be forthcoming as experience is gained in handling real data. The simulators will continue to be useful tools in ensuing development stages.

The remarks made in Section 4, particularly items a, b, and d, will be expanded upon below as recommendations for possible consideration.

Because of the uncertainty in timing and the infrequent rate of the LTN-51 data which furnish the SICLOPS System with heading and track angle and velocity, we suggest that three similar and available outputs from the Litton navigation system be incorporated in the ADAS recording system. These three outputs are true heading, which is available in the Litton analog output as synchro data, and the north and east components of velocity which are available as binary outputs. All three of these outputs are updated every 50 msec by the Litton system with a timing uncertainty of 12 msec. The increased data rate of the heading information would then provide the SICLOPS System with adequate yaw data. The binary outputs (the two velocity components) would have resolutions of 0.1 knot as compared to a resolution of 1 knot in the current LTN-51 velocity record. Executing this change would eliminate the timing uncertainty and also reduce the data preparation effort since the LTN-51 record would no longer be needed. The smoothing technique to reduce the effects of noise, as now used on current ADAS data, could then be applied to these three new inputs.

Another suggestion is to relocate the LTN platform to the immediate neighborhood of the scanner on-board the aircraft. The dynamic deformation of the fuselage of the aircraft in flight (possibly amounting to several degrees) could misalign the Litton unit which is mounted aft, relative to the scanner which is mounted forward. Surveying the LTN platform to the scanner platform would be very desirable.

BIBLIOGRAPHY

- Bendix Corporation, Multispectral Scanner System (Operation and Maintenance Instructions-Preliminary Technical Manual), MSDS-MSS-1, Bendix Corp., Aerospace Systems Division, Ann Arbor, December 1970, (Prepared for NASA).
- Department of the Army, Universal Transverse Mercator Grid, Technical Manual No. TM 5-241-8, Dept. of the Army, Washington, 7 July 1958.
- Earth Resources Data Format Control Book, PHO-TR543, Volume 1, March 1973, Philco-Ford Corp., Houston Operations, for Flight Support Division, NASA/MSC.
- Fails, Anthony M., A Detailed Guide to the Processing of Earth Observation Aircraft Program Guidance and Navigation Data, Report No. LEC-0206, February 1973, Lockheed Electronics Co., Inc., Houston Aerospace Systems Division, for Computation & Analysis Division, NASA/MSC.
- Hasell, P. G. and L. M. Larsen, Calibration of an Airborne Multispectral Optical Sensor, Willow Run Laboratories, Institute of Science and Technology, The University of Michigan, Ann Arbor, September 1968, U. S. Army Electronics Command Technical Report ECOM-00013-137.
- Lambeck, P. F., Engineering Evaluation of 24-Channel Multispectral Scanner, Report No. WRL 10822-1-T, Willow Run Laboratories, Institute of Science and Technology, The University of Michigan, Ann Arbor, (Prepared for NASA/MSC), 1972.
- Litton System, Inc., Litton LTN-51 Inertial Navigation System (Installation Training Manual, Document No. TT404, Aero Products Division, Litton Systems, Inc., Woodland Hills, Cal., April 1971.
- Litton Systems, Inc., Litton LTN-51 Inertial Navigation System (With Dual DME Update Option) for Commercial Aviation (Product Description), Publication No. 6242, Aero Products Division, Litton Systems, Inc., Woodland Hills, Cal., August 1971.
- Litton Systems, Inc., Litton LTN-51 Inertial Navigation System for Commercial Aviation (Technical Description), Document No. 633001, Aero Products Division, Litton Systems, Inc., Woodland Hills, Cal., December 1972.
- McClure, Connie L., Theory of Inertial Guidance, Prentice-Hall, Englewood Cliffs, N. J., 1960.
- National Aeronautics and Space Administration, Aircraft Remote Sensing Systems, MSC-05118, NASA Manned Spacecraft Center, Houston, December 1971, (Prepared by Lockheed Electronics Company, Inc.).
- Program Documentation, Earth Observation LTN-51 Navigation Data Processing Program (LTNDAT), Program Q656, Project 4699 F, DPD#0060, August 1970, Lockheed Electronics Company, Houston Aerospace Systems, for Computation & Analysis Division, NASA/MSC.
- Program Documentation, Guidance & Navigation Rate Program for Earth Resources (ASQRAT), Program Q636, Project 4699, DPD#0064, August 1970, Lockheed Electronics Company, Houston Aerospace Systems, Houston, for Computation & Analysis Division, NASA/MSC.
- U. S. Air Force, Controlled Range Network, RC-010903, Recon Division, AF Avionics Laboratory, Air Force Systems Command, Wright-Patterson Air Force Base, Ohio, June 1965 (Maintained and Operated by Data Corp., Dayton, Ohio).

Appendix AXIS SYSTEMS

The positive direction sense for the three Euler angles of the B-frame (aircraft) to the E-frame is as follows: Heading is nose right, pitch is nose up, and roll is right wing down.

For the projection of the xy plane of the E-frame onto flat terrain, the x and y axes follow the same convention as the E-frame (i.e., x points North and y points East). The magnitude of the ground velocity is V; the direction is given by the track angle, β , measured clockwise from North.

Body Axes (B-Frame)

Conventional body axes x_b , y_b and z_b are used with origin at the vehicle's center of gravity. In level flight the axes for a right-hand set have the x_b axis pointing forward, the z_b axis downward, and y_b axis sideward. For a winged vehicle the x_b , z_b axes lie in the plane of symmetry of the vehicle.

Euler Angles (E-Frame)

Orientation of the vehicle body axes is specified by three conventional Euler angles: ψ (heading), θ (pitch), and ϕ (roll); these give the orientation of B-frame with respect to E-frame. The E-frame origin is the center of gravity of the vehicle. The reference plane is perpendicular to a radius vector to the center of the earth; the reference direction in this reference plane is North. Thus, for a spherical earth, the unit vectors \bar{i}_e , \bar{j}_e and \bar{k}_e for the E-frame point North, East and downward, respectively.